

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</small>					
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)

Concurrent Security of Cryptographic Protocols:
FA9550-08-1-0197
Final Report

Rafael Pass
Department of Computer Science
Cornell University
Ithaca, NY 14853

January 11, 2012

1 Introduction

Following the ground-breaking work on public-key encryption in the 70's, the field of Cryptography has evolved far beyond securing message transmission. Today, cryptographic protocols are used in large-scale systems to guarantee not only confidentiality and authenticity but also attack- and fault-tolerance. For instance, the notion of a *secure computation*, introduced by Yao and Goldreich, Micali and Wigderson in the early 80's, enables a set of parties to, through the execution of a distributed communication protocol, securely implement any service that a trusted party could perform for them. Security here means that, even if an arbitrary subset of the parties get corrupted and deviate from their prescribed instructions, both correctness and confidentiality is still maintained.

This novel use of cryptography, however, also admits new types of attacks. The security of most cryptographic protocols (and, in particular, those for secure computation) can be compromised if many instances of the protocol are concurrently executed. This concurrent setting allows a coordinated attack in which an adversary controls many parties, interleaving the executions of the various protocol instances. For instance, a so called *man-in-the-middle* attacker participating in two simultaneous executions of a cryptographic protocol might use messages from one of the executions in order to violate the security of the second.

Consider a two-party protocol between A acting as an *initiator*, and B acting as a *responder*. A man-in-the-middle adversary M controlling the channel between A and B can participate in an interaction with A , acting as a responder, and at the same time participate in an interaction with B , acting as an initiator. Furthermore, by exploiting the interaction with A , M might be able to violate the security of the interaction with B . At a first glance, it seems that such an attack can be prevented by encrypting all communication between A and B . This does not work: If M is acting as truthful responder in its interaction with A , then A will believe that M is the rightful owner of the messages she sends, and thus encrypt all her messages using M 's key. The same holds for B . Indeed Lowe's famous attack on the Needham-Schroeder protocol works this way.

On the Internet concurrent attacks are unavoidable. While both the need and definitions were articulated in the early 90's, constructions of concurrently secure protocols are lacking. During the reporting period, we have developed several novel techniques for dealing with concurrent attacks, leading to the resolution of several decade-old open problems:

- We obtained the first sub-logarithmic construction for defending against man-in-the-middle attacks based the minimal assumption of one-way functions; this had remained a major open problem for almost 20 years.
- We obtained the first secure computation protocol that relies on the same assumptions as the original work by Goldreich, Micali and Wigderson but uses a sublinear number of communications rounds, resolving a problem open since the conception of secure multi-party

computation in 1987.

- We demonstrated how techniques used to defend against man-in-the-middle attacks can be leveraged to get concurrent secure multi-party computation protocols; this connection made is possible to significantly weaken earlier assumptions on trusted infrastructure and/or computational intractability assumptions.
- We constructed the first secure computation protocols that require no trusted infrastructure other than authenticated communication, and that satisfy a meaningful notion of security that is preserved under concurrent executions assuming standard cryptographic hardness assumptions.
- We demonstrated several new techniques for constructing practical and concurrently secure protocols for the specific class of, so-called, zero-knowledge protocols.

As outlined in our original proposal, our work has focused on two interrelated threads:

1. **Minimizing Trusted Set-Up**
2. **Defending Against Man-in-the-middle Attacks**

Below, we elaborate on some of our most central contributions in these two threads, and conclude with a brief summary of some other results that were obtained as a consequence of a joint study of the above threads.

2 Minimizing Trusted Set-up

To avoid the complexities concurrency brings, most of the literature relies on a *trusted set-up* during which all participating parties are initialized with trusted information. Typically, a trusted set-up requires an initial intervention of a trusted party; execution of the protocol can then proceed without further help from that trusted party. But in many settings it might be hard to agree on a trusted party or such a party might not exist.¹

An alternative method for obtaining a trusted set-up is to assemble all participating parties in a “trusted isolated chamber” where no communication with the outside world is possible. The parties by themselves then run a secure computation protocol to provide the appropriate set-up. After this initialization phase, they can subsequently engage in executions of cryptographic protocols—with the same group of parties—but this time at distance. This approach is not practical because it requires assembling all parties at the same physical location.

¹The situation is exacerbated by the fact that in known solutions the trusted party can cheat in an arbitrary and undetectable way.

Using AFOSR we made significant progress on minimizing the trusted set-ups required for executing concurrently secure protocols (see publications [23,15]). In particular, we showed that significantly weaker (and thus harder to compromise) set-up suffices. To give one example, the most popular trusted set-up model is the common reference string (CRS) model, introduced by Blum, Feldman and Micali, which assumes (1) all parties participating in the protocol execution have access to a common string that has been “ideally” sampled from some specific pre-defined distribution, and (2) that no side information about how the string was generated is known to any protocol participant. Protocols in the CRS model are typically quite particular about the distributions they need. Some protocols specify distributions that involve some non-trivial sampling process, but even in protocols that need relatively simple distributions (say, the uniform distribution) the security analysis quickly falls apart as soon as the distribution of the common reference string is changed only slightly. The sensitivity of the CRS model to distribution peculiarities rules out physical implementations where the reference string is taken to be the result of joint measurement of some physical phenomenon. While it is reasonable to believe that such phenomena are largely unpredictable and uncontrollable, so they have high entropy, it is hard to believe they are taken from an exact distribution that is known to and useful for the protocol designer. As a corollary of our general result, we establish that indeed, as long as the common reference string has sufficient amount of *min-entropy*² concurrent security can be established.

Perhaps more importantly, our methodology provides a general framework for constructing concurrently secure protocol. This approach unifies, simplifies and improves, essentially all known results on concurrent secure computation. The key idea behind this framework is a general method for leveraging techniques used to defend against man-in-the-middle attacks (see Section 3) to get full concurrent security.

Although our initial goal was to simply lower the assumptions on the trusted set-up, our framework leads to improvements also in efficiency and computational assumptions. Perhaps surprisingly, it even improves known results without concurrency: For instance, the work presents the first asymptotic improvement to the round-complexity of the seminal general secure computation protocol of Goldreich, Micali and Wigderson without making any further assumptions; this had remained open since 1987.

3 Defending Against Man-in-the-middle Attacks

One of the most basic—and easiest to mount—attacks on a cryptographic protocol, involves a man-in-the-middle adversary having full control over the communication channel linking the parties executing the protocol. *Non-malleable* protocols defend against man-in-the-middle attacks. The design and analysis of non-malleable protocols (i.e., protocols that withstand man-in-the-

²The *min-entropy* of a distribution D is $-\log_2 p$, where p denotes the probability of the most likely string in D . Thus, roughly speaking, a source with min-entropy k , has “as much” randomness as a k -bit random string.

middle attacks) is a notoriously difficult task. The task becomes even more challenging if honest parties do not use any kind of trusted set-up. Indeed, only a handful of such protocols have been constructed so far.

We here focus on one of the most basic and central cryptographic tasks: commitment schemes. Often described as the digital analog of sealed envelopes, *commitment schemes* enable a party, known as the *sender*, to commit itself to a value while keeping the value secret from the *receiver*. Furthermore, the commitment is *binding*, meaning that in a later stage when the commitment is opened, a single value determined in the committing stage will be found. The application, ranges from coin flipping to secure computation.

Non-malleable commitments get to the essence of a man-in-the-middle adversary. For example, an adversary upon seeing a commitment to a specific value v , might be able to commit to a related value (say, $v - 1$), even though it does not know the actual value of v . Non-malleable commitments block such attacks.

The original paper by Dolev, Dwork and Naor (DDN) presents the first protocols achieving non-malleable commitments without relying on any trusted set-up. Their protocols only require the existence of one-way functions (which is the minimal cryptographic assumptions) but require $O(\log n)$ rounds of interaction, where $n \in N$ denotes the length of a party identifier (e.g., an IP-address). Since their seminal work over 2 decades ago, there has been a large literature attempting to improve the round-complexity of non-malleable commitments. But all such attempts require either trusted set-up assumptions, or stronger computational intractability assumptions; no improvements to the round complexity of non-malleable commitment assuming minimal assumptions were found.

During the reporting period, we managed to address this central open problem, and go beyond the DDN $O(\log n)$ -round barrier for constructing non-malleable commitments based on minimal assumptions: assuming only one-way functions, we construct an $O(1)^{\log^* n}$ round protocol.³ See publication [14]. Rather than providing a direct construction of a full-fledged non-malleable commitment, we start with a weak notion of non-malleability; we, next, present a technique for *amplifying* weakly non-malleable commitments to full-fledged ones. The amplification procedure consists of two basic steps that are iterated $\log^* n$ times. Since the original publication of this work, several other works have relied on this “non-malleability” amplification procedure.

Other methods for defending against man-in-the-middle attacks are explored in publications [17, 19, 3].

³Recall that \log^* (the iterated logarithm function) is the number of times the logarithm function must be iteratively applied before the result is less than or equal to 1. For instance $\log^* 2^{65536} = \log^* 2^{2^{2^2}} = 5$.

4 Other Significant Results

- **Practical Methods for Defending Against Concurrent Attacks.** Our most general techniques for defending against concurrent attacks are still inefficient. We have also focused on developing practical solutions for certain specific types of cryptographic protocols. One central task of interest is so-called *zero-knowledge proofs*: Zero-knowledge proofs (introduced by Goldwasser, Micali and Rackoff) are protocols that enable one party—called the *prover*—to convince another party—called the *verifier*—about the validity of some mathematical statement without revealing anything else about the content of the statement. Such protocols are useful in preventing active attacks on cryptographic protocols. Specifically, adversarial behavior is prevented by requiring protocol participants to prove in zero-knowledge that they have followed the protocol. As such, zero-knowledge proof also serve as a test-bed for the more complex notion of secure computation. In publications [20,18,16, 10,7,6] we have developed several new techniques for improving the efficiency of zero-knowledge protocols, and establishing practical solutions to concurrently secure zero-knowledge protocols. We have also developed several novel lower-bounds demonstrating what type of inefficiencies are necessary.
- **Concurrent Security Without Any Trusted Set-up.** By combining techniques developed in the above two threads, we constructed the first secure computation protocol that does not require *any* trusted infrastructure other than authenticated communication, and that satisfies a meaningful notion of security that is preserved under concurrent executions assuming standard cryptographic hardness assumptions; see publication [2]. This addresses a central problem left open since the seminal works on concurrent security from 1990. Such new models of protocols security are further explored under the auspices of our AFOSR grant entitled “New Models of Protocols Security”
- **Game-theoretic approaches to cryptographic protocols.** During the reporting period we also initialized a game-theoretic study of cryptographic protocol security; see publications [13,9]. This research direction is further explored under the auspices of our AFOSR grant entitled “New Models of Protocols Security”.

5 List of Publications

1. T. Roeder, R. Pass and F. Schneider. Multi-Verifier Signatures. To appear in *Journal of Cryptology*, 2010.
2. R. Canetti, H. Lin and R. Pass. *Adaptive Hardness and Composable Security from Standard Assumptions*. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2010)*, pages 541-550, 2010. Invited to *SIAM Journal of Computing*, special issue on selected papers of FOCS 2010.

3. H. Lin, R. Pass, W. Tseng and M. Venkatasubramanian. Concurrent Non-Malleable Zero Knowledge Proofs. *Advances in Cryptology (CRYPTO 2010)*, Springer LNCS 6223, pages 429–446, 2010.
4. R. Pass and H. Wee. Constant-round Non-malleable Committments from Subexponential One-way Functions. *Advances in Cryptology (EUROCRYPT 2010)*, Springer LNCS 6110, pages 638–655, 2010.
5. J. Halpern and R. Pass. I Don’t Want to Think about it Now: Decision Theory with Costly Computation. *Proceeding of the 12th International Conference on the Principles of Knowledge Representation and Reasoning (KR 2010)*, 2010.
6. R. Pass, M. Venkatasubramanian and W. Tseng. Eye for an Eye: Efficient Concurrent Zero Knowledge in the Timing Model. In *Proceedings of the 7th Theory of Cryptography Conference (TCC 2010)*, pages 518–534, 2010.
7. R. Pass and M. Venkatasubramanian. On Public versus Private Coins in Zero-Knowledge Proofs. In *Proceedings of the 7th Theory of Cryptography Conference (TCC 2010)*, pages 588–605, 2010.
8. R. Pass, J. Hastad, D. Wikstrom and K. Pietrzak. An Efficient Parallel Repetition Theorem. In *Proceedings of the 7th Theory of Cryptography Conference (TCC 2010)*, pages 1–18, 2010.
9. J. Halpern and R. Pass. Game Theory with Costly Computation: Formulation and Application to Protocol Security. In *Proceeding of the 1st Innovations in Computer Science Conference (ICS 2010)*, 2010.
10. R. Pass, W. Tseng and D. Wikstrom. On the Composition of Public-coin Zero Knowledge. In *Advances in Cryptology (CRYPTO 2009)*, Springer LNCS 5677, pages 160–176, 2009. Full version to appear in *SIAM Journal of Computing*, 2011.
11. J. Halpern and R. Pass. Iterated Regret Minimization: A New Solution Concept. In *Proceeding of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 153–158, 2009. Full version to appear in *Games and Economic Behavior*, 2011.
12. J. Halpern and R. Pass. A Logical Characterization of Iterated Admissability. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2009)*, pages 146–155, 2009.
13. J. Halpern, R. Pass and V. Raman. An Epistemic Characterization of Zero Knowledge. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2009)*, pages 156–165, 2009.
14. H. Lin and R. Pass. Non-malleability Amplification. In *Proceedings of the 41th Annual Symposium on Theory of Computing (STOC 2009)*, pages 189–198, 2009.

15. H. Lin, R. Pass and M. Venkatasubramanian. A Unified Framework for Concurrent Security: Universal Composability from Stand-alone Non malleability. In *Proceedings of the 41th Annual Symposium on Theory of Computing (STOC 2009)*, pages 179–188, 2009.
16. R. Pass and H. Wee. Black-box Constructions of Two-Party Protocols from One-way Functions. In *Proceedings of the 6th Theory of Cryptography Conference (TCC 2009)*, pages 403–418, 2009.
17. O. Pandey, R. Pass and V. Vaikuntanathan. Adaptive One-Way Functions and Applications. *Advances in Cryptology (CRYPTO 2008)*, Springer LNCS 5157, pages 57-074, 2003.
18. R. Pass and M. Venkatasubramanian. On Constant-Round Concurrent Zero Knowledge. *Proceedings of 5th Theory of Cryptography Conference (TCC 2008)*, pages 553–570, 2008. Invited to Journal of Cryptology.
19. H. Lin, R. Pass and M. Venkatasubramanian. Concurrent Non-malleable Commitments from One-way Functions. *Proceedings of 5th Theory of Cryptography Conference (TCC 2008)*, pages 571–588, 2008. Invited to Journal of Cryptology.
20. O. Pandey, R. Pass, A. Sahai, W. Tseng and M. Venkatasubramanian. Precise Concurrent Zero Knowledge. *Advances in Cryptology (EUROCRYPT 2008)*, Springer LNCS 4965, pages 397–414, 2008.
21. R. Pass, A. Shelat and V. Vaikuntanathan. Relations Among Notions of Non-malleability for Encryption. *Advances in Cryptology (ASIACRYPT 2007)*, Springer LNCS, pages 519–525, 2008.
22. R. Cramer, G. Hanaoka, D. Hofheinz, H. Imai, E. Kiltz, R. Pass, A. Shelat and V. Vaikuntanathan. Bounded-CCA Secure Encryption. *Advances in Cryptology (ASIACRYPT 2007)*. Springer LNCS, pages 502–518, 2008.
23. R. Canetti, R. Pass and A. Shelat. Cryptography from Sunspots: How to Use an Imperfect Reference String. *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007)*, pages 249–263, 2007.